Case Study: Teaching Research Skills to Computer Science Graduate Students

Marie desJardins
Department of Computer Science and Electrical Engineering
University of Maryland Baltimore County
Baltimore, MD 21250
mariedj@cs.umbc.edu

Abstract

The UMBC Department of Computer Science and Electrical Engineering has developed a new class on "Basic Research Skills for Computer Science" for Ph.D. and M.S. students.[1] The class was first offered in Spring 2004, and will be taught for the second time in Spring 2005. The class focuses on teaching five basic skills that are necessary for research: critically reading and evaluating technical papers, performing literature surveys, developing research ideas, writing technical papers and proposals, and giving oral presentations. In this paper, the structure of the class is described. Quantitative and qualitative data is presented to support our claim that the course benefits students and teaches them important skills.

INTRODUCTION

In Fall 2002 and Spring 2003, the UMBC Computer Science faculty had a series of discussions about the best way to prepare students to become effective researchers. The faculty believed that in addition to specific technical knowledge, general research skills were important for student success. We identified a set of five such skills: (1) critically reading and evaluating technical papers, (2) performing literature surveys, (3) developing research ideas, (4) writing technical papers and proposals, and (5) giving oral presentations. Although these skills were often gained by individual students over the course of the program—from the student's advisor, by observation, or from peers—there was no systematic mechanism in the curriculum for teaching them directly to students. The committee felt that teaching these skills explicitly, early in a student's graduate career, would better prepare them for the thesis or dissertation process and for a research career.

The result of these discussions was a proposal and syllabus for a new class on "Basic Research Skills for Computer Scientists," designed for M.S. and Ph.D. students in the first year of the program. The objectives of the course were (1) to give students a clear understanding of the steps to successfully complete the degree, and (2) to give them a solid grounding in the five research skills listed above. One of the major open questions in designing the course was whether it was possible to teach these skills without a specific technical focus. The proposed course attempted to address this by requiring each student to choose a specific research area within their own discipline, and by requiring an outside reader for the major written components of the class. Judging from student feedback, it appears that the course did successfully teach general research skills within this framework, and that students viewed the course as a very positive experience.

I taught the first offering of the course in Spring 2004 to a class of 13 M.S. and Ph.D. students. This paper describes the course, discusses its success at achieving the stated objectives, and presents our plans for future course offerings. The course was taught for the second time in Spring 2005, and is expected to become a regular course in the CS graduate curriculum at UMBC.

COURSE OVERVIEW

The course is structured around a "mini-research project," and includes a series of written exercises, lectures and discussions on writing skills, oral presentations by students, lectures and discussions on other research skills, and guest speakers addressing topics relevant to graduate students. These course components are discussed in more detail below.

Preliminaries. The first class session gives an overview of the course, covers basic LaTeX commands, and includes an extensive discussion of plagiarism issues and proper citation styles. Two technical papers from the same research area are assigned for the next class discussion. Students read these papers and write paper summaries (see Figure 1) in preparation for the class. During the discussion, we analyze and critique the two papers in depth, for technical content, writing style, and overall presentation. We also discuss how submitted papers are reviewed by conference and journal reviewers.

Research Project. The core of the research class is a structured, supervised "mini-research project." In the first part of the semester, each student performs a literature survey on a topic of their choice. In the second part of the

[1] Course materials for the Spring 2004 offering are available at http://www.csee.umbc.edu/~mariedj/691b/ . Course materials for the Spring 2005 offering are available at http://www.csee.umbc.edu/~nicholas/691B/ .

Guidelines for Paper Summaries (shortened from original handout)
Paper summaries should not be longer than two pages, and should generally be one page or less. You want to hit all of the key points, but not get into too much detail, or it's not really a *summary*.
Summaries should provide answers to these questions about the paper:

- What is the specific problem that the work is intended to address?
- What is the authors' main claim? (i.e., what are they trying to do?)
- What evidence do the authors give to support their claim? Is the evidence adequate?
- Does the work appear to be original and significant?
- Do the authors place their work in the context of related work? Ideally, they should cite relevant literature and should explain why their work represents an original and significant contribution.
- Is the paper well organized?
- Is the paper clearly written?

Some other questions you may wish to think about and answer, in order to make your paper summary as thorough as possible:

- What assumptions or limitations does the work have?
- What would be the logical next steps for the research?
- What are your remaining open questions about the work after reading the paper?
- What are the most important citations in the paper to follow up on if you were to read further on this topic? Why?

Strategies for Reading Papers: If you've come across a paper in the course of your explorations, you need to first decide if it's relevant. (Obviously, if it's an assigned paper, you may skip this step!)

- Decide from the title and context if it might be relevant.
- If so, read the abstract, then read through the introduction and conclusion quickly.
- Once you've decided to read the paper, don't just dive into it. Approach it strategically.
  - Skim the whole paper once to get a feeling for what it's about. Skip anything dense or technical. Glance at the figures and tables.
  - Read through the whole paper, but don't let yourself get stuck. If you have questions, jot a note in the margin and keep going.
  - Go back and work through the details of any (important) equations, proofs, algorithms, etc. that you skipped over.
- If you're going to present the paper at a lab meeting or class, or if it's very relevant to your own research, you may need to spend more time with the paper to fully understand it. You may also need to look up citations of work on which this paper builds.
- Even for papers that you don't need to turn in a research summary on, It's a good idea to write a short summary of the paper in your research journal or on the first page of the paper.

Figure 1: Guidelines for Research Summaries

semester, each student develops a research proposal on their chosen topic, and writes a formal paper summarizing the related work and proposed research. Each student is required to identify an outside reader (i.e., a faculty member other than the instructor) who agrees to read and provide feedback on the literature survey and the research paper.

Written Exercises. In addition to the research project, a number of written exercises are included. These include writing short paper summaries for several technical papers discussed during class; developing a research portfolio, including a CV; and developing a professional website. Class attendance and participation accounts for a significant part of the grade.

In order to encourage good writing style through feedback and iteration, students are required to turn in multiple versions of the two major papers. For the literature survey, they must submit an annotated bibliography, an outline of the literature survey, a draft survey, and a final survey. For the research paper, they must submit a research problem summary, a proposal draft, a full paper draft, and the final paper. Students must also produce detailed written reviews of two other students' draft papers, along the lines of a conference paper review.

Writing Skills. The course includes several lectures and discussions about writing issues: paper organization, general writing style, preparing bibliographies, common errors, and useful tools. We spend one class doing an in-depth analysis of two survey papers, and talking about guidelines for preparing good literature surveys (see Figure 2).

Finding Relevant Papers:

- Google is a useful tool, but should not be the only place you look!
- CiteSeer, an online citation index and paper database maintained by NEC, is a terrific resource, as is Google Scholar.
- Identify a few important and relevant papers, and work forwards and backwards through citation links.
- It's important to know what are the key publications (top journals and conferences) and researchers (most published and cited authors) in your field of interest. Ask somebody who's knowledgeable.
- If you see a citation or a name repeatedly, look it up! It's important! But also note: "important" is not synonymous with "good" – sometimes everybody cites a paper just because everybody else does.
- Pay attention to institutions – you'll quickly learn which places are doing important work in your area.

Locating and Reading Papers:

- If you can't find a paper online, try the library, other students, your advisor or outside reader.
- Don't make the mistake of "depth-first search" of the paper space. Instead, look at papers closely enough to know how important they are, and to start creating clusters of "similar" papers (i.e., identifying themes, threads, or styles of research within the field). Then organize your reading by clusters. This is a much more efficient way to read than a scattershot approach. Also, doing this as you go along will help you to organize the literature survey itself.
- As you read papers, make note of important citations to follow up on (and what cluster those citations seem to belong to). Know when to stop. There will be more related fields than you can possibly become an expert in.
- Take notes and keep them in an organized system – notebook, online, whatever. Don't just scribble on pieces of paper. I like to make notes in the margins of papers, but also to write short bulleted summaries of papers that are particularly relevant.

Figure 2: Literature Survey Guidelines

Oral Presentations. Each student is required to select a key paper from their literature survey and present a summary of the paper. Later in the semester, students must give an oral presentation of their research proposal. Every student is asked to fill out a detailed feedback form for each of the other students' presentations. This process not only gives the students feedback on their own talks, but also gives them greater insight into what works (and what doesn't work) in the other students' talks.

Other Research Skills. Additional lecture topics include how to organize and present talks effectively, finding a good research topic, finding an advisor, developing a research idea, empirical methodologies in CS, statistical methods, preparing research proposals, and grant writing. Assigned readings for these parts of the class are from Zobel (1997) and Peters (1997).

Guest Speakers. In Spring 2004, an invited speaker from NIH/NLM also talked about proposal writing from a funder's perspective, and also gave some tips on interviewing for academic and research positions. To break up the series of student presentations, one day's class consisted of a career panel, with four invited speakers giving their views on industry, research laboratory, government, and academic careers.

STUDENT CHARACTERISTICS

The Spring 2004 class consisted of 13 students with diverse research interests and personal backgrounds.

Diversity of Research Interests. Most of the students in the class were planning to do research in artificial intelligence or graphics. One student was focused on computer security, and one on computer arithmetic. My primary research area is artificial intelligence; I knew virtually nothing about graphics, security, or computer arithmetic. Similarly, many of the students had very little background in some of the areas in which the other students were doing research.

I was initially quite concerned that these different interests would lead to a lack of communication within the class, and in particular that I would have difficulty providing substantive feedback on written assignments in research areas with which I was unfamiliar. In fact, I found nearly the opposite. Because of my relative lack of familiarity with some of the subject areas, it was clear in those papers when there was missing background, logical inconsistencies, or fundamental flaws in a proposed approach. To make their points understood to a non-expert, the students were forced to write (and often rewrite) in a very clear, direct style. On the other hand, because of my familiarity with the AI topics, it was sometimes more difficult for me to notice the gaps or flaws in reasoning in those papers: my knowledge of the area would subconsciously tend to fill any holes in the presentation.

The diversity of student interests also enabled the students to give each other a wide range of feedback on their oral presentations. Students in similar areas would point out specific problems with the descriptions of existing work, or suggest missing references. Students in other areas would comment on missing background material or suggest the need for clearer examples to illustrate complex ideas.

#	Question	1	2	3	4	5	Mean	Dept	Campus	Rank
1	Did you gain new insights/skills?	0	1	1	2	9	4.46	4.15	4.23	534/1470
2	Did the instructor make clear the expected goals?	0	0	2	3	8	4.46	4.15	4.23	534/1470
3	Did other evaluations reflect the expected goals?	0	0	3	3	6	4.25	4.19	4.24	766/1413
4	Did assigned readings contribute to what you learned?	0	0	4	1	8	4.31	3.59	4.02	481/1391
5	Did written assignments contribute to what you learned?	0	0	0	4	9	4.69	4.03	4.10	164/1347
6	Was the grading system clearly explained?	1	0	1	1	10	4.46	4.19	4.16	492/1455
7	How would you grade overall teaching effectiveness?	0	0	2	1	7	4.50	3.99	4.08	332/1451
8	Were the instructor's lectures well prepared?	0	0	0	7	5	4.42	4.36	4.43	875/1407
9	Did the instructor seem interested in the subject?	0	1	0	0	11	4.75	4.65	4.70	788/1404
10	Was lecture material presented and explained clearly?	0	1	0	2	8	4.55	4.14	4.28	511/1401
11	Did the lectures contribute to what you learned?	0	0	2	1	9	4.58	4.15	4.27	520/1398
12	Did A/V techniques enhance your understanding?	0	0	1	5	6	4.42	4.06	3.96	334/1176
13	Did class discussions contribute to what you learned?	1	0	0	1	7	4.44	3.87	4.06	395/1276
14	Were all students actively encouraged to participate?	0	0	0	0	9	5.00	4.03	4.30	1/1277
15	Did the instructor encourage fair and open discussion?	0	1	0	2	6	4.44	3.94	4.28	623/1271

Figure 3: **SCEQ Data for CMSC 691B, Spring 2004.** The columns labeled 1 through 5 indicate the number of students who gave each rating for that question. "Mean" indicates the mean rating for that question for the course. "Dept" and "Campus" indicate the mean rating for that question all courses taught within the department and at UMBC, respectively, in Spring 2004. "Rank" is the rank of the class for that question of all classes taught at UMBC in Spring 2004. Light gray rows are those for which the course was ranked above the campus mean. Dark gray rows are those for which the course was ranked at or above the top campus-wide quartile.

Diversity of Backgrounds. The national origins of the students in the class included the U.S. (both white and African-American students), China, India, and Poland. Two of the 13 students were female. Again, rather than the cultural and linguistic differences presenting barriers, they proved to be opportunities. The students learned to generate written and oral presentation material that could be understood by a broad audience, and learned to appreciate and handle the linguistic challenges that can arise when trying to communicate with people with different backgrounds and accents.

COURSE EVALUATION

The evaluation of the course is based on UMBC's Student Course Evaluation Questionnaire (SCEQ) and on written comments from the students. The SCEQ, which is administered in all classes at UMBC, asks each student to rate the course quantitatively in a number of different areas. The written comments come from the "blue sheet" component of the SCEQ, which allows students to provide additional anonymous feedback to the instructor, and from e-mails sent by individual students.[2]

Although the small sample size ($n = 13$) obviously means that strong conclusions cannot be drawn about the effectiveness of the class, the overall very positive response from students can be taken as an indication that the class may have significant benefits for students.

[2]Only selected questions that appeared most relevant are included in Table 3. Omitted questions include how often class was cancelled, questions about exams (which were not given), and specific questions relevant primarily or only to laboratory classes, seminar classes, and field work.

Quantitative SCEQ Data

Figure 3 shows the quantitative SCEQ data for the course. Rows highlighted in light gray indicate evaluation criteria for which the course was ranked above the campus mean (all but one question; a sign test for this result gives a confidence of $p = 0.0034$ that the course is above average). Rows highlighted in darker gray indicate evaluation criteria for which the course was ranked at or above the top quartile (about a third of the questions).[3] The ratings clearly show that the students generally had a favorable impression of the course. The only criterion for which the course rating was below the campus mean was Question 8, about lectures being well prepared, for which the class was only 0.01 below the campus mean. (Since this was a new course, being slightly below the mean for "well preparedness" is not surprising.)

Overall teaching effectiveness (Question 7) was rated 4.5 (compared to a department mean of 3.99 and a campus mean of 4.08), with a campus-wide ranking of 332 out of 1451 classes evaluated for this question. That the course was generally beneficial is also supported by the response to Question 1: most students felt that they gained new insights and skills from the class.

[3]UMBC is a medium-size state university that emphasizes high-quality teaching. As such, student satisfaction with teaching is fairly high, so the campus mean is a competitive benchmark. (Note that the campus means for most questions are at or above 4 out of 5.) Also, as at many institutions, computer science teaching quality is generally perceived to be lower than campus-wide teaching quality (note the consistently lower marks for department-wide means, as compared to campus-wide means). Therefore, for a CS class to outperform the campus mean indicates a strong positive statement about student perceptions of the class.

Indirect evidence that the course objective of increasing the students' ability to read and evaluate technical papers is given by the high rating for Question 4; evidence of improving written skills by Question 5; and evidence of improving oral communication skills by Questions 13 and 14.

Qualitative Student Feedback

One month into the class, 8 of the 9 students who responded to an informal (anonymous) survey ranked the class a 5 (on a scale from 1–5) for the question "Of the classes you've taken or are taking at UMBC, how useful would you say this class is?"

Dr. Jack Prostko, the Director of Faculty Development at UMBC, came to the class halfway through the semester to perform a small-group evaluation. (I was not in the room during the evaluation.) The students were asked to discuss what was working well in the class, what was not working well, and how to improve the class.

The discussion below summarizes comments from the small-group feedback session and the SCEQ blue sheets.

General Student Feedback. In general, the students had very positive things to say about the class. The students seemed to particularly appreciate becoming more familiar with the research cycle, learning to read papers quickly and effectively, gaining confidence in their written and oral presentation skills, and receiving useful, detailed feedback about their writing. Non-native speakers of English found it especially useful that they were consistently pushed to write and speak extensively in English, which they did not always have the opportunity to do in their other classes.

Specific student comments from the blue sheets:[4]

- Thanks a lot for the course! I enjoyed it a lot!!
- The best part of the course was learning to skim papers quickly, writing a lot of papers/summaries and gaining confidence in writing as well as presenting.
- I learnt a lot about graduate research, expectations both of and from students. Also paper writing (proposal, grant, etc.) were very important part of the learning experience.
- [The best part of the course was] class open discussion on readings; gave option to see how others viewed topic.
- Excellent course, must be made compulsory for all graduate students.
- [The instructor] spent *a lot* of time reading and commenting on written assignments. Great amount of structure especially considering it is the first time the course has been offered.
- [The instructor's] interest in her work and teaching made the class more exciting and informative.
- [The best part of the course was] the insight into graduate school from a professor's standpoint, lectures on funding and proposals, writing and reviewing the survey/proposal.
- It gives us so much info about graduate school and specifically computer science. I learned a lot of useful skills.
- [The instructor was] very passionate about teaching this course. The instructor is very open and brings a lot of personal experience into the teaching of this course. This is exactly what I am expecting.
- [The best part of the course was] frequent feedback, possibility to focus on research (and some pressure to do this in a timely fashion).
- This course gives us a lot of research skills such as writing summary, writing survey paper, writing proposal, doing a good presentation. All these is very helpful for students who don't do research before.

Specific Suggestions for Improvement. The students did have several suggestions for improving the course. To summarize and respond briefly:

1. *Provide less closely spaced deadlines.* This would be difficult to do without dropping some of the assignments altogether, although it might be possible to move some of the deadlines earlier, to encourage students to start thinking about the research paper sooner.
2. *Assign less work.* I did, in fact, drop a couple of the lesser written assignments, in an attempt to reduce the workload.
3. *Bring more of the other faculty members in to talk about their research or to guest-lecture on selected topics.* I did try to do this, but the response from the faculty was limited. Possibly in future semesters, given the positive response of students to these activities, the faculty will be more motivated to participate.

Some of the assignments (a library exercise, being required to meet with faculty members other than the advisor) were not perceived to have much value. Also, some of the students disliked being required to read papers in other disciplines, or felt that the viewpoint was too limited to AI and graphics. Typical comments were "the course should be more related to direct research" and "research methodology needed more specific perspectives." To address this issue, some students suggested splitting some of the lectures/discussions into smaller groups by discipline, for more focused discussion in specific research areas and in discipline-specific methodologies.

Specific quotes from student blue sheets:

- I would like to recommend the organization of course to focus more on the research portfolio and spend more time on working on writing final paper.
- It was hectic at times due to all the other courses. Might want to consider offering it in the summer.
- Should also be asked to read other students MS/PhD dissertations since we all have to do it at some point.
- [Drop] the "Writing for Computer Science" book.
- I don't think we should be required to meet with 3 professors unless we do not have an advisor.

[4] All student quotes are direct and unedited except where elided in square brackets.

- Would have been nice to have more on proposal budgets.
- Possibly [have the course] co-taught by faculty members in different departments *only* for discussions such as experimental methodology. Break the class up into groups for the discipline-specifi c parts and keep them together for statistical discussions. More experimentation discussions.
- Many more instructors experienced in research should present, instead of one teacher.

Additional Student Feedback. From a student who was not registered for the class, but sat in for part of the semester:

This mail is for nothing but to thank you for offering 691B.

Thank you very much for admitting me to your Basic Research Skills class. Although I was unable to fully attend, the benefits of the class to me have been plenty. For your information, please let me list them:

- *I have read many useful resources on your homepage (on the schedule page). They are really [worthwhile]. I wish I had enrolled to a similar class at my first semester. I have saved each of those resources, and will review them very often.*
- *Although I had some Latex tradition from undergraduate years, I have rediscovered many aspects of it recently.*
- *I have quickly prepared a bibtex bibliography and progressed fast in managing references, which was difficult before.*

I strongly believe that those who have taken this class are making a jump-start into research. I hope the department offers similar courses in the future, as well.

Excerpts from another message, from a student who did take the class:

I've been meaning to write to you for a long time about a few things. Firstly, I enjoyed the 'Basic research skills' course a lot and I wish you were teaching a course related to computer graphics in some way, so that I could take your course :)

Anyway, about CMSC691b I wanted to suggest that you should please try and make it a requirement for ALL PhD students and those MS students who might be interested in pursuing a PhD in the long run. I think it was an extremely beneficial course and it helps me approach research more professionally.

At the same time, I wanted to suggest to you that you write a book on the lines of the course. I completely understand that you are busy with your research and teaching, and would be willing to help out in any and every way to see such a beneficial book come out in print. Please feel free to drop me a mail whenever you need any assistance with such an endeavor.

Please let me know if you think about writing the book or any other ideas that might improve our department. I could team up with a few students and attempt at implementing it.

CONCLUSIONS AND FUTURE PLANS

Overall, both the students and I felt that the course experience was a very positive one. I enjoyed teaching it (much more than I thought I would), and the students seemed to genuinely appreciate the opportunity to focus on developing basic skills in a supportive environment. The positive outcomes that have resulted for the students are perhaps best summed up by an e-mail I received from a professor in the Biology department, writing about a Ph.D. student who took the class in his second semester at UMBC:

I am just writing to let you know that I am very impressed with the impact that your "Research Skills for Graduates" course has had on B—, the CS student with whom I am working.

Specifically, I recently received my first formal research 'write-up' from B—, and was so impressed with the clarity and overall quality with which he wrote (quite beyond the technicalities of content) that I asked him whether this was a natural strength. He told me that, quite the reverse, it was a skill he had learned through much hard work in your course. So "Thank you!" — I will probably be seeking to institute some similar kind of course for our biologists.

The class was taught for the second time in Spring 2005, with 18 students registered. We plan to follow up on this paper with additional data and suggestions for faculty teaching similar classes at other institutions.

Acknowledgements

Thanks to Penny Rheingans and Scott Cost, the other members of the subcommittee that originally developed the course proposal, for their inputs and ideas. Thanks also to Tom Dietterich (Oregon State University) and Susan Epstein (City University of New York) for sharing their experiences and course materials from related courses, and to Dragos Margineantu (OSU / Boeing) for sharing his experience as a student in Prof. Dietterich's class.

References

Peters, R. L. 1997. *Getting What You Came For: The Smart Student's Guide to Earning a Master's or Ph.D. (Revised Edition)*. Farrar, Straus and Giroux.

Zobel, J. 1997. *Writing for Computer Science*. Springer.